THE COERULEAN CLOCK

The Timing of Locus Coeruleus Loss in Alzheimer's Disease — Why the First Neurons to Sicken Are Among the Last to Die, How the Axon Dies Decades Before the Cell Body, and What Occupies the Interval Between Them

Three Clocks, Not One · The Third Decade · The Dying Back
The Late Death · The Second Trigger · The Interval of Sleep

A Doctoral Synthesis Prepared in Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy in Neuroscience

Prepared under the Organic Network Synthesis methodology

AdultCognitiveDisease.com

July 2026

The chronology volume of the coerulean series. The locus coeruleus is the first structure in the brain to accumulate tau and among the most depopulated at the end. Reciting those two facts together has concealed the interval between them — some twenty-five years in which the nucleus is sick and not dying. This volume dates that interval, and finds the axon being removed inside it while every cell body is still present and countable.

"The long gap between NFT accumulation and neuronal loss suggests that a second trigger may be necessary to induce neuronal death in AD."

— Theofilas and colleagues, *Alzheimer's & Dementia*, 2017

This volume is an attempt to answer that sentence. It takes the gap seriously as a duration rather than a delay, asks what is happening inside it, and finds that the answer is not nothing — it is the quiet removal of an arbor, one axon at a time, from cells that a pathologist would still count as alive.

THE COERULEAN SERIES THE CHRONOLOGY VOLUME

The First Ember (how tau is kindled in the locus coeruleus: DOPEGAL, AEP, the guardian withheld)

The Coerulean Shears (the protease and the hand: norepinephrine drives MMP-9 against the sulfated surface)

The Coerulean Pincer (the dual-armed assault converging on one tau kinase)

The Silenced Eraser (the ungoverned microglion and the silencing of PP2A)

The Coerulean Clock ← this volume

The four preceding coerulean volumes are volumes of mechanism. They establish how tau is kindled in this nucleus, what its deranged output does to the matrix outside the target neuron and to the kinase within it, and how its failure releases the microglion. None of them asks the question that governs whether any of it can be reached: *when*. This volume asks it, and finds that the nucleus does not degenerate on one schedule but on three — a tau clock starting in the third decade, an axon clock that finishes before the cortex has a single tangle, and a somatic clock that does not begin until a third of the nucleus is already gone. The clinically decisive one is the middle clock, which no routine measurement in the field is designed to see.

Contents

Abstract

The locus coeruleus is the first structure in the human brain to accumulate abnormal tau, and it is among the most severely depopulated structures in the end-stage Alzheimer brain. These two facts are almost always recited together, and reciting them together has produced a persistent and consequential error: the assumption that they describe one process, running continuously from its beginning to its end. They do not. Between the appearance of tau in a coerulean neuron and the disappearance of that neuron there lies an interval of decades, and this dissertation is about what happens inside it.

The evidence for the interval is quantitative and it is not in dispute. Hyperphosphorylated tau inclusions are present in 7.9% of locus coeruleus neurons at Braak stage 0 — that is, in brains with no cortical tangle pathology whatsoever — and pretangle material is demonstrable in the nucleus from the second and third decades of life (Ehrenberg and colleagues, 2017; Braak and Del Tredici, 2011). Yet unbiased stereology of the same nucleus across the full staging series finds that neuronal number is essentially preserved through Braak stages 0–II and begins to fall only from stage III onward, so that "neuronal loss started only midway through AD progression" (Theofilas and colleagues, 2017). The gap between the two events has been estimated at a quarter of a century (Matchett and colleagues, 2021). Theofilas and colleagues drew the correct inference and stated it as an open question: the long gap between tangle accumulation and neuronal loss implies that tau alone does not kill the coerulean neuron, and that **a second trigger** must be required. This dissertation is an attempt to say what that second trigger is, and to date the events on either side of it.

The central claim is that the locus coeruleus does not degenerate on one clock but on three, which run in a fixed order and are separated by decades. **Clock one is the tau clock**: it starts in the third decade, in the cell body, and it runs almost silently. **Clock two is the axon clock**: it starts far earlier than the field has assumed, and it is the clock that matters clinically. In the *App*$^{NL\text{-}G\text{-}F}$ mouse, noradrenergic fibre loss in a coerulean target field is already 14% at two months, 27% at three and 33% at six, while the number of locus coeruleus cell bodies remains statistically indistinguishable from wild-type at twelve months; in human tissue the corresponding fibre degeneration is present in Braak stage 1–2 material and does not worsen thereafter (Meyer and colleagues, 2025). The axon is not lost because the cell body died. The axon is lost while the cell body is alive, healthy by count, and firing. **Clock three is the somatic clock**: it starts at Braak III and runs to a loss of roughly 8.4% of coerulean volume per Braak stage, through a 30% neuron deficit already present at amnestic mild cognitive impairment and a further 25% by mild-to-moderate dementia, to an end-stage figure of 63–83% (Theofilas 2017; Kelly and colleagues, 2017; Zarow and colleagues, 2003). By the time the somatic clock is audible, the disease is over as a therapeutic proposition. The clinical action is entirely inside the interval, on clock two.

We then ask what runs clock two, and propose an answer assembled from five literatures that do not cite one another. The coerulean neuron is an autonomous pacemaker with a vast, thin, poorly myelinated arbor — a single such neuron accompanies some 20 metres of cerebral capillary — and it is constitutionally without a perineuronal net, so that the reelin-staged brake on its tau kinase is never assembled in the first place (Matchett 2021; Morawski and colleagues, 2010). In the diseased state these neurons fire *more*, not less; and the axonal consequence of firing more is calcium entry, phosphatidylserine externalisation on the axolemma, recognition by the bridging molecule MFG-E8, and phagocytic removal of the axon by microglia that never touched the cell body (Meyer 2025). Firing, in other words, is what eats the axon. This reframes the second trigger as a *failure of excitability control*, and identifies a specific and testable brake: **KCNIP4**, the Kv4-channel-interacting protein that shapes the A-type potassium current, which has now been shown to be upregulated early and selectively in the neurons that resist Alzheimer's disease, and whose overexpression suppresses the activity markers Arc and c-Fos (Dharshini and colleagues, 2026). A neighbouring wake-promoting nucleus supplies the proof of principle: hypocretin neurons become hyperexcitable in ageing through loss of a different potassium brake, and disrupting that brake in young animals reproduces the aged phenotype (Li and colleagues, 2022).

Four further molecules are placed in the chronology, each at the station where the evidence puts it and no earlier. **PININ** — the nuclear-speckle protein encoded by *PNN*, whose abbreviation collides unhappily with the perineuronal net and which is a different object entirely — is depleted from nuclear speckles by tau aggregates along with SRRM2, an event accompanied by the retention of more than 1,200 introns across 641 genes (Lester and colleagues, 2021). This places a splicing lesion *inside* the coerulean cell body during the silent interval, and supplies a candidate mechanism by which the tau clock could set the pace of the excitability clock: a neuron whose splicing apparatus is disassembled cannot reliably build the isoform-dependent brake that *KCNIP4* encodes. We grade this bridge as inference, not evidence. **C4d**, the durable covalent split product of complement C4, has been shown to bind LilrB2 with nanomolar affinity, to colocalise with it at excitatory synapses in human cortex, and to mediate synapse pruning — and to rise with age and further in Alzheimer's disease (Brott and colleagues, 2025). It is the tag that marks the synapse for removal at the far end of the coerulean axon. **EPHEXIN-5**, the RhoA exchange factor that acts as a developmental brake on excitatory synapse formation and is normally degraded by EphB2 signalling, is elevated in the hippocampi of human Alzheimer's patients, is induced by amyloid-beta, and when genetically reduced rescues spine density and cognition in *hAPP* mice (Margolis and colleagues, 2010; Sell and colleagues, 2017). It is the ratchet: it does not remove the synapse, it forbids its replacement, converting a reversible loss into a permanent one. **REELIN** enters twice — as the brake the locus coeruleus never had, and as the marker of the neurons that receive the coerulean lesion next, since reelin-immunoreactive layer II entorhinal neurons are precisely the population that first accumulates intracellular amyloid in early disease (Kobro-Flatmoen and colleagues, 2016).

Finally we place sleep, which is not an epilogue to this argument but its tempo control. The locus coeruleus is the one cell population in the brain that is granted a scheduled and total off-period: its neurons fire at roughly 2 Hz in waking, fall to under 1 Hz in slow-wave sleep, and cease almost entirely in REM (Aston-Jones and Bloom, 1981). If firing is what condemns the axon, then REM is the axon's reprieve, and its loss is not a symptom but a lesion. NREM makes the complementary demand: the infra-slow noradrenergic oscillations of NREM sleep are the pump that drives glymphatic clearance, and abolishing them abolishes the clearance (Hauglund and colleagues, 2025). The locus coeruleus is therefore required to oscillate in one sleep state and to fall silent in the other, and a sickening nucleus fails at both — a second instance of the arsonist operating the sprinkler. In cognitively unimpaired older adults, worse coerulean structural integrity already tracks with more frequent nocturnal awakenings (Van Egroo and colleagues, 2021).

We close with a graded validity ledger separating what is measured from what is inferred, a statement of what would falsify the three-clock model, and the therapeutic consequence that follows from it: that the target is not the coerulean cell body, which is not dying when the disease can still be reached, but the coerulean axon, which is — and that the two interventions the model nominates are an excitability brake and an intact night.

I. The Question Is When, Not Whether

That the locus coeruleus degenerates in Alzheimer's disease is among the best-established facts in the neuropathology of the condition, and it has been established for forty years. Bondareff and his colleagues reported it in 1982, finding a loss of some 80% of coerulean neurons in a young, severely demented subgroup; German and his colleagues reconstructed the whole nucleus in 1992 and mapped the topography of the loss; Zarow and his colleagues, comparing subcortical nuclei head to head in eighty-six pathologically confirmed cases, found the coerulean loss to be the most severe of all — 83% in Alzheimer's disease, exceeding the loss in the nucleus basalis and in the substantia nigra pars compacta, and exceeding the loss in the substantia nigra even in Parkinson's disease. The nucleus is small: some 98,000 neurons in a structure of about 13 cubic millimetres. Its depopulation is not subtle, and no one disputes it.

The question this dissertation asks is a different one, and it is the question the field has been slowest to answer with numbers rather than adjectives. *When* does the locus coeruleus die? Not whether, not how much, but on what schedule — and, critically, whether the schedule of the cell body is the same as the schedule of everything that cell body owns.

The reason this matters is not academic. Nearly every therapeutic proposal that has been advanced for the noradrenergic system in Alzheimer's disease implicitly assumes a particular answer. Proposals to replace norepinephrine assume that the deficit is one of supply and that the supply fails when the cells die. Proposals to protect coerulean neurons assume that the cells are dying at the moment the patient becomes reachable. Proposals to image the locus coeruleus as an early biomarker assume that what the image measures — signal intensity, contrast ratio, estimated integrity — is tracking the same process at the beginning of the disease that it tracks at the end. Each of these assumptions is a claim about timing, and each of them is, on the evidence assembled here, wrong in a specific and correctable way.

The single-clock picture and why it fails. The received account runs as a smooth curve. Tau appears in the locus coeruleus very early — earlier than anywhere else in the brain. The nucleus then degenerates progressively, losing cells across the decades, until by the end of the disease most of it is gone. Norepinephrine falls with the cell count; the cortex is deprived; cognition suffers. It is a single monotonic process with an early start and a late finish, and the natural inference is that the tau one sees at the beginning is the cause of the loss one counts at the end, acting slowly but continuously in between.

Two observations break this picture, and they break it in opposite directions.

The first is that the early tau does not kill. In the stereological series assembled by Theofilas and colleagues, spanning the whole Braak range in human post-mortem brainstem, coerulean neuronal number was preserved through Braak stages 0, I and II, and declined only from stage III. Their own summary is unambiguous: "neuronal loss started only midway through AD progression." Yet by Ehrenberg's count, 7.9% of coerulean neurons already carry a hyperphosphorylated tau cytoplasmic inclusion at Braak stage 0 — in brains, that is, with no cortical tangles at all — and that proportion has doubled by stage I. Tau is present, in a substantial minority of the nucleus, for a period during which the nucleus loses no measurable number of cells. Matchett and colleagues put the separation at "at least 25 years." Whatever tau is doing to the coerulean neuron in that quarter-century, it is not killing it.

The second observation is stranger, and it is the one this dissertation is organised around. During the same interval in which the cell bodies are demonstrably surviving, the *axons* are demonstrably not. In the *App*$^{NL\text{-}G\text{-}F}$ knock-in mouse, Meyer and colleagues measured noradrenergic fibre density in the olfactory bulb and found a 14% deficit at two months of age, 27% at three months, and 33% at six months. They then counted the locus coeruleus neurons that those fibres came from, at twelve months — twice the age of the last fibre measurement — and found no difference from wild-type. In human tissue they saw the same dissociation: pronounced degeneration of norepinephrine-transporter-positive fibres in olfactory bulbs from prodromal and early Alzheimer cases at Braak stage 1–2, with no further decline in more advanced cases. A third of the projection was gone while the projecting population was intact.

Put these two observations together and the single-clock picture collapses. What one is looking at is not one process with a long time-course but at least three processes with different onsets, different rates, and different terminations — and the one that is easiest to measure, the loss of cell bodies, is the last of the three to begin and therefore the least useful for anything one might wish to do about it.

Three clocks. The organising proposal of this dissertation can be stated in a sentence. The locus coeruleus in Alzheimer's disease runs three clocks: a **tau clock** in the soma, which starts in the second-to-third decade and runs for twenty-five years without killing anything; an **axon clock** in the projection, which starts early enough to be complete in the target field by Braak stage I–II and which strips the neuron of its function while sparing its life; and a **somatic clock**, which starts at Braak III, is well advanced by the time a patient reaches a memory clinic with amnestic mild cognitive impairment, and finishes with the near-obliteration of the nucleus. They are not three stages of one process. They are three processes, with different mechanisms, and only the middle one is both early enough and consequential enough to be worth attacking.

The dissertation proceeds by dating each clock in turn (Sections II–IV), assembling them into a single chronological table (Section V), and then asking what drives the middle clock — which is

where the interesting biology, and all six of the molecules named in this volume's brief, turn out to live.

A note on what "loss" means, and on the trap in the imaging literature. Because this dissertation is about timing, it must be pedantic about what each method actually measures, and one distinction deserves stating before any evidence is cited. Unbiased stereology on post-mortem tissue counts *cell bodies*, and counts them well. Neuromelanin-sensitive magnetic resonance imaging *in vivo* measures a signal that depends on neuromelanin content, on cell size, on water content and on the geometry of the nucleus, and it is conventionally reported as "locus coeruleus integrity" — a word that quietly elides the distinction this dissertation exists to draw. A coerulean neuron that has lost a third of its axonal arbor, shrunk its soma, contracted its dendrites and downregulated tyrosine hydroxylase is a neuron with degraded "integrity" and an intact entry in the stereological count. The imaging literature is therefore probably more sensitive to the early interval than the counting literature is, and the frequent framing of the two as measuring the same thing at different resolutions is an error. When Bueichekú and colleagues report that coerulean integrity changes *precede* medial temporal tau accumulation, and Jacobs and colleagues report that *in vivo* integrity tracks neuropathology and cognitive decline, they are not contradicting Theofilas's finding that cells do not die until Braak III. They are measuring the interval that Theofilas's method, by design, cannot see.

II. Clock One — The Tau

The tau clock is the best-dated event in the natural history of Alzheimer's disease, and it is dated earlier than any other.

The third decade. Braak and Del Tredici examined brains from individuals under thirty years of age and found abnormal, non-argyrophilic tau — pretangle material, not yet the silver-positive filamentous tangle — in the locus coeruleus. The finding was extended in the large series that gives the modern staging its foundation: 2,332 unselected brains from individuals aged one to one hundred, examined with AT8 immunocytochemistry and Gallyas silver staining, in which the earliest tau lesions in the entire brain were coerulean and appeared in the young. The nucleus is not merely the first cortical-projecting structure to be affected. On this evidence it is affected in people who will not develop dementia for fifty years, and in a substantial number who never will.

Braak's own account distinguishes a graded series of pretangle stages — the "a, b, c" of pretangle tau — that precede the mature neurofibrillary tangle and that begin in the young. The material starts as a soluble, abnormally phosphorylated species distributed through the somatodendritic compartment, and only later condenses into the argyrophilic inclusion that the classical stains detect. This distinction matters for timing because the classical Braak staging scheme, which is defined on argyrophilic tangles in the cerebral cortex, is by construction blind to the coerulean pretangle. Braak stage 0 does not mean "no tau." It means "no cortical tangles," and the locus coeruleus at Braak stage 0 is already, in Ehrenberg's count, 7.9% affected.

The quantitative anchor. Ehrenberg and colleagues applied unbiased stereology to forty-eight well-characterised cases enriched for controls and early stages, counting hyperphosphorylated tau neuronal cytoplasmic inclusions in sixty-micron sections of locus coeruleus and dorsal raphe. Their figure for Braak stage 0 — 7.9% of coerulean neurons bearing an inclusion, against 2.6% in the dorsal raphe — is the cleanest available statement of how early the coerulean lesion is, because it is anchored to a stage at which the rest of the brain is, by definition, clean. Matchett and colleagues, reviewing the same literature, note that the proportion doubles by Braak stage I and reaches essentially the entire nucleus — 100% of neurons containing tau pathology — by Braak stage VI.

That last figure is worth pausing on. At the end of the disease every surviving coerulean neuron contains tau. And yet the nucleus is not empty: 17–37% of its neurons are still there, depending on the series. Tau presence is therefore neither sufficient for death — the neurons carry it for decades and survive — nor, in the terminal state, discriminating, since it labels the survivors as thoroughly as

it labelled the departed. Whatever selects which coerulean neurons die and when, it is not the mere presence of tau in them.

Why here first. The question of why the locus coeruleus is the brain's first tau lesion has been answered, in the literature, by an accumulation of partial reasons rather than a single one, and the accumulation is itself the point: the nucleus is not vulnerable for one reason but for six, and they compound.

Its axons are long, thin, and poorly or incompletely myelinated, which raises the metabolic cost of conduction, raises oxidative burden, and leaves an enormous membrane area exposed to whatever the extracellular space contains. Its arbor is extraordinary in extent — Matchett and colleagues note that each coerulean neuron accompanies, on average, some twenty metres of cerebral capillary, a coverage exceeding that of any other neuronal type. It is an autonomous pacemaker, firing without needing excitatory drive, and paying for that autonomy in continuous activity-dependent calcium entry and the mitochondrial oxidative load that follows. It sits against the fourth ventricle and innervates the great majority of the brain's microvasculature, positions that maximise its exposure to whatever arrives from the circulation and the cerebrospinal fluid. It manufactures norepinephrine, whose metabolite DOPEGAL activates asparagine endopeptidase, which cleaves tau at N368 into an aggregation-prone fragment — an autotoxicity available only to a catecholaminergic cell, and the subject of the companion volume *The First Ember* (Kang and colleagues, 2020). And, finally, it is one of the subcortical nuclei that Morawski and colleagues identified as constitutionally devoid of an aggrecan-based perineuronal net, in a survey whose central finding was that net-ensheathed neurons resist the tangle and net-less nuclei do not.

That last item is where reelin enters the chronology, and it enters as an absence. The reelin brake on glycogen-synthase-kinase-3-beta requires N-sulfated heparan sulfate as an obligate co-receptor, presented on a sulfated extracellular surface; the perineuronal net is that surface; the locus coeruleus does not have one. The companion volumes of this series have argued at length that the reelin brake is *cut* in the cortex by proteolysis. In the locus coeruleus it is not cut. It was never assembled. The nucleus enters life without the protection that the rest of the brain spends the disease losing, which is a satisfying explanation for why it is affected first and — as Section XI takes up — a discouraging one for anyone hoping to protect it by restoring reelin signalling.

What the tau clock does not do. The decisive negative result about clock one is Theofilas's, and it has been stated already but deserves restating in its own terms because everything downstream depends on it. Across the full Braak series, coerulean volume falls by 8.4% per stage increment — a smooth, early-starting, monotonic decline. Coerulean *neuronal number* does not: it holds through stages 0 to II and falls from III. Volume and number therefore dissociate in exactly the interval this dissertation is about. Something is shrinking the nucleus for decades before anything is killing it — and the candidates for what shrinks a nucleus whose cells are all still present

are somatic atrophy, dendritic contraction, and the loss of the neuropil that the axons and dendrites constitute. Matchett and colleagues describe precisely this: contracted dendrites and swollen cell bodies as the early morphological signature.

The tau clock, then, produces a sick, shrunken, dendritically contracted, still-living neuron, and it produces it decades before anything dies. It is a clock of morbidity, not mortality. Theofilas and colleagues, confronting the same arithmetic, wrote the sentence that this dissertation takes as its brief: *the long gap between NFT accumulation and neuronal loss suggests that a second trigger may be necessary to induce neuronal death in AD.*

III. Clock Two — The Axon

The second clock is the one the field has systematically underweighted, for a reason that is methodological rather than conceptual: cell bodies are easy to count and axons are not. A coerulean neuron is a soma of perhaps twenty microns attached to an arbor of metres. Standard neuropathology measures the twenty microns.

The dissociation, measured. The clearest existing measurement of the axon clock comes from a target field rather than from the cortex at large, and the choice of target field is itself instructive. Meyer and colleagues (2025) studied the olfactory bulb, on the reasoning that olfactory dysfunction is among the earliest clinical signs in Alzheimer's disease and that the bulb receives a dense noradrenergic innervation from the locus coeruleus. In the *App*$^{NL\text{-}G\text{-}F}$ knock-in mouse they quantified norepinephrine-transporter-positive fibres and found progressive loss: 14% at two months, 27% at three months, 33% at six months. They then asked the question that makes the study decisive. They counted the locus coeruleus neurons themselves, at twelve months — six months after the last fibre timepoint, at an age by which any dying-back process driven by somatic death should have shown itself — and found no difference in neuron number between *App*$^{NL\text{-}G\text{-}F}$ and wild-type animals.

A third of the projection to one target had been removed from a nucleus that had not lost a single countable cell.

The human arm of the same study is the more important half, because mouse models of amyloid deposition are not the disease. In human olfactory bulb tissue from prodromal and early Alzheimer cases — Braak stage 1–2, Thal phase 1–2 — the authors found pronounced degeneration of norepinephrine-transporter-positive fibres relative to age-matched unaffected controls. And then, crucially, the loss "did not further decline in progressive AD cases." The axon clock in this target field runs to completion before Braak stage III — before, that is, the somatic clock has started. Two clocks, running consecutively rather than concurrently, with the axon clock finishing where the soma clock begins.

One should not over-generalise from the olfactory bulb, and Section XIII grades this appropriately: the bulb is a single target field, and the claim that the whole coerulean arbor follows the same schedule is an extrapolation. But the extrapolation is supported from an independent direction. Matchett and colleagues report that coerulean loss is not uniform along the nucleus but sharply rostral-predominant — an 83% loss of length in the rostral portion against 23% in the middle and 15% caudally — and rostral coerulean neurons are precisely those projecting to the hippocampus, entorhinal cortex and neocortex. Dahl and colleagues, imaging *in vivo*, found rostral

coerulean integrity to be the portion associated with memory performance in older adults. A degeneration organised by projection target rather than by cell position within the nucleus is what one expects if the lesion begins at the terminal.

Why the axon dies first: the mechanism. The mechanism Meyer and colleagues identify is the part of their work most relevant here, because it supplies a candidate for Theofilas's second trigger, and because it is not the mechanism the field would have predicted.

The axon is not starved and it does not simply wither. It is *eaten*. Olfactory bulb microglia in *App*$^{NL\text{-}G\text{-}F}$ mice showed a 33% higher phagocytic capacity within twelve hours, and the axons were recognised by them through phosphatidylserine externalisation on the axonal membrane — the canonical "eat-me" signal — bridged to the microglial phagocytic apparatus by MFG-E8. Genetically reducing phagocytosis preserved both the axons and the olfaction. The removal of the coerulean axon is therefore an immune act, executed by a phagocyte on a living neuron's process, and it is reversible in principle because it is regulated rather than degenerative.

And what causes the phosphatidylserine to appear on the axon? The authors' answer is the hinge of this dissertation. Locus coeruleus neurons in the *App*$^{NL\text{-}G\text{-}F}$ mice showed an overall increase in spontaneous action potential frequency, and the externalisation of phosphatidylserine was calcium-dependent. The neuron's own firing floods the axon with calcium; the calcium flips phosphatidylserine to the outer leaflet; the microglion reads the flag and removes the process.

The coerulean axon is destroyed by the coerulean neuron's own activity.

This is a genuinely different account of neurodegeneration from the one the field usually offers, and its consequences run through the rest of this volume. It means that the sick coerulean neuron is not a passive victim. It means that the compensatory hyperactivity documented in the human nucleus — the raised tyrosine hydroxylase message in surviving neurons, the dendritic sprouting into the peri-coerulean zone, the axonal sprouting to the hippocampus that Szot and colleagues reported in 2006, the rising cerebrospinal-fluid norepinephrine that climbs with disease severity (Elrod and colleagues, 1997) — is not merely a compensation that fails. It is a compensation that *costs*, and the currency it is paid in is axon. And it means that the therapeutic instinct to stimulate the failing locus coeruleus, which has considerable experimental support at the level of cognition, carries a specific and previously unnamed risk at the level of the axon.

It also, incidentally, explains a puzzle in the human data. If firing drives axonal removal, and if surviving neurons fire harder to compensate for those already lost, then the axon clock should be self-accelerating in its early phase and should decelerate once enough somata have gone that compensation collapses. Meyer's human finding — pronounced fibre loss at Braak 1–2 that does not further decline in advanced cases — is exactly that shape.

A negative result that matters: complement is not the tag on the axon. Because this dissertation must also place C4d, it is worth recording precisely what Meyer and colleagues found when they looked for complement. They examined C1q decoration of the norepinephrine-transporter-positive axons and found *no* significant change in C1q colocalisation between genotypes. Complement is not the signal that marks the coerulean axon for removal; phosphatidylserine and MFG-E8 are.

This is a useful discipline. The complement-tagging story is so well established for the synapse in Alzheimer's disease that it is tempting to extend it to every act of microglial removal, and here that extension is specifically excluded by direct measurement. C4d has a place in this chronology — Section VIII gives it one, and it is a well-evidenced place — but the place is the synapse, not the axon. Two different removals, two different tags, at two different points on the same neuron.

IV. Clock Three — The Soma

The third clock is the one the literature has always measured, and the numbers are consistent across four decades and several methods. What has been missing is not the data but its placement in time relative to the other two clocks.

Onset at Braak III. Theofilas and colleagues' stereological series is the reference measurement because it spans the whole staging range rather than contrasting two endpoints. Its two results should be held apart. Coerulean *volume* declines by 8.4% for each one-unit increase in Braak stage — a decline that begins immediately and runs throughout. Coerulean *neuronal number* is preserved through stages 0–II and declines only from stage III. The authors also report that age-related change spares the locus coeruleus, which removes the obvious confound: this is disease, not ageing, and the nucleus is not one of those structures that thins in everybody.

The clinical translation. Kelly and colleagues (2017) approached the same nucleus through clinical rather than pathological staging, applying tyrosine-hydroxylase immunohistochemistry and unbiased stereology to cases characterised as no cognitive impairment, amnestic mild cognitive impairment, or mild-to-moderate Alzheimer's disease. They found a 30% loss of coerulean neurons already in the transition from no impairment to amnestic mild cognitive impairment, and a further 25% loss in the progression to Alzheimer's disease — approximately a 50% deficit in the dementia group relative to unimpaired controls.

At first reading this appears to contradict Theofilas. It does not. Amnestic mild cognitive impairment corresponds, in the great majority of cases, to Braak stages III–IV. Kelly's "loss by MCI" and Theofilas's "loss begins at stage III" are the same finding expressed in two staging languages, and the convergence is a good deal more reassuring than either study alone. The translation is worth stating explicitly because it is the hinge of the clinical argument: **by the time a patient presents with amnestic mild cognitive impairment, roughly a third of the locus coeruleus is already gone.** The somatic clock does not begin at diagnosis. It is a third of the way through.

Kelly's molecular arm makes the same point in a different currency. Single-population microarray profiling of the surviving coerulean neurons showed significant reductions in functional classes of messages governing mitochondrial respiration, redox homeostasis and neuritic structural plasticity — cytochrome C1, glutathione peroxidase, neurofilament subunits — in both the amnestic mild cognitive impairment and the Alzheimer groups, and the reductions tracked global cognitive deterioration and neuropathological burden. Note the third class. *Neuritic structural plasticity*: the surviving neurons have downregulated the machinery of building and maintaining processes.

This is the somatic clock's molecular signature and it is also, read from the other side, a statement about the axon — the neuron has stopped investing in the arbor that Section III showed being removed.

The endpoint. Zarow and colleagues' comparison across subcortical nuclei gives the terminal figure and the comparative context: 83% coerulean neuronal loss in Alzheimer's disease, 68% in Parkinson's disease, in both cases exceeding the loss in the nucleus basalis and the substantia nigra pars compacta. Matchett and colleagues cite an average of 63% across the literature. The spread between 63% and 83% reflects case selection and method rather than genuine disagreement; the qualitative statement is not in doubt, which is that the end-stage locus coeruleus has lost most of itself.

Oh and colleagues (2019) add a discriminating observation. Studying three wake-promoting nuclei — the locus coeruleus among them — across Alzheimer's disease, progressive supranuclear palsy and corticobasal degeneration, they found that all three tauopathies accumulate considerable tau in these nuclei and show a decrease in neurotransmitter-synthesising neurons, but that *substantial neuronal loss was exclusively found in Alzheimer's disease*. Tau in a wake-promoting nucleus is common to the tauopathies. Death of that nucleus is specific to Alzheimer's disease. This is the strongest available argument that the somatic clock is driven by something other than tau burden alone — which is, once again, Theofilas's second trigger, arriving from a third direction.

Which neurons die. The somatic clock does not run at the same rate throughout the nucleus, and its topography is informative. Matchett and colleagues report an 82% decrease in large multipolar neurons against a 39% decrease in small fusiform neurons, and the rostrocaudal gradient already noted — 83% length loss rostrally, 23% in the middle, 15% caudally. Larger neurons with larger arbors, projecting to the regions that fail first, die preferentially. If the axon clock is the primary lesion, this is exactly the expected pattern: the neurons that lose the most axon are the neurons with the most axon to lose.

What the somatic clock costs. Wilson and colleagues (2013), following 165 participants of the Rush Memory and Aging Project through a mean of 5.8 years of annual cognitive testing to autopsy, found that the density of noradrenergic neurons in the locus coeruleus was associated with the rate of cognitive decline, and concluded that coerulean neuronal density may be a structural component of neural reserve. The nucleus is therefore not merely an early casualty; its remaining population is a determinant of how fast everything else goes. This is the bridge to Section XII, and it is a two-way bridge: the locus coeruleus is both the first thing the disease damages and one of the things that determines how much damage the rest of the brain sustains.

V. The Chronological Summary

The three clocks can now be placed on one axis. The table below is the analytical core of this dissertation, and every row is sourced. Ages are approximate and are given for a sporadic case reaching dementia in the eighth decade; the ordering, not the absolute dating, is what the evidence supports.

Approx. age	Braak	Locus coeruleus — soma	Locus coeruleus — axon & terminal	Clinical state	Key evidence
20–30 yr	Pre-0	Pretangle tau appears; not yet silver-positive; somatodendritic. No loss.	No measured change.	Silent	Braak & Del Tredici 2011; Braak et al. 2011
30–50 yr	0	7.9% of neurons bear p-tau inclusions. Number intact.	Not measured in humans at this stage.	Silent	Ehrenberg et al. 2017
~50 yr	I	Tau-bearing fraction doubles. Number still intact. Volume falling ~8.4%/stage.	**Pronounced NET+ fibre degeneration already present in target field.**	Silent to subtle (olfaction)	Theofilas et al. 2017; Meyer et al. 2025
~55–65 yr	II	Number intact. Somatic atrophy, dendritic contraction.	Fibre loss established; does not further decline afterwards.	Subjective decline; sleep fragmentation	Matchett et al. 2021; Van Egroo et al. 2021
~65–70 yr	III	**Somatic clock starts.** Neuronal loss first detectable.	Compensatory sprouting of surviving axons to hippocampus.	Amnestic MCI (~30% LC neurons already lost)	Theofilas et al. 2017; Kelly et al. 2017; Szot et al. 2006
~70–75 yr	IV	Loss continues at ~8.4% volume/stage; large multipolar neurons preferentially.	CSF norepinephrine rising with severity — surviving output deranged upward.	Mild dementia	Kelly et al. 2017; Elrod et al. 1997
~75–80 yr	V	Cumulative loss approaching 50%+ vs. unimpaired.	Cortical noradrenergic innervation grossly deficient.	Moderate dementia	Kelly et al. 2017
80+ yr	VI	63–83% loss. 100% of survivors contain tau. Rostral 83% length loss.	Arbor largely gone.	Severe dementia	Zarow et al. 2003; Matchett et al. 2021

Three features of this table carry the argument.

First, the ordering is fixed and the gaps are enormous. Tau precedes somatic death by a period estimated at twenty-five years or more. That is not a delay to be explained away; it is the longest and most accessible therapeutic window in the entire natural history of the disease, and it is currently unoccupied by any intervention.

Second, the axon and the soma are on different clocks and the axon runs first. The row that carries the most weight is Braak stage I, where a target field shows pronounced fibre degeneration at a stage where cell counts are normal. Functional deafferentation of coerulean targets is therefore an *early*-stage event, and coerulean cell death is a *mid-to-late*-stage event, and treating them as the same lesion has caused the field to look for the disease's noradrenergic contribution in the wrong decade.

Third, output and number move in opposite directions in the middle of the table. Between Braak III and V the nucleus is losing neurons while the norepinephrine reaching the cerebrospinal fluid is rising. This is the "rising tide" documented in the companion volume *The Coerulean Pincer* — compensatory hyperactivity of survivors, raised tyrosine hydroxylase message, dendritic and axonal sprouting — and Section VI now supplies the reason it is not merely a curiosity but the engine of the second clock.

VI. The Second Trigger — Excitability and the Brake That Fails

Theofilas and colleagues asked what second trigger converts a tau-bearing but living coerulean neuron into a dying one. This section proposes that the trigger is not a new insult arriving from outside but the failure of a control system inside — specifically, the failure of the neuron's own excitability brake — and that the resulting hyperactivity is what runs clock two.

The coerulean neuron is a pacemaker, and pacing is expensive. Unlike most cortical neurons, locus coeruleus neurons are autonomously active: they fire tonically without requiring excitatory synaptic drive, at roughly one to three hertz in the waking animal. That autonomy is maintained by activity-dependent calcium entry, and Matchett and colleagues list it among the primary reasons for the nucleus's selective vulnerability, because continuous calcium cycling imposes a continuous mitochondrial oxidative load. A neuron of this design has no idle state. Its baseline is a load.

The observed derangement is upward. Three independent lines converge on the conclusion that the sickening coerulean system does not simply fall silent. Szot and colleagues found, in post-mortem Alzheimer and dementia-with-Lewy-bodies tissue, three changes in the surviving noradrenergic neurons all consistent with compensation: increased tyrosine hydroxylase message, sprouting of dendrites into the peri-coerulean dendritic zone, and sprouting of axonal projections to the hippocampus. Elrod and colleagues found cerebrospinal-fluid norepinephrine rising with the severity of the disease. And Meyer and colleagues, recording directly, found an overall increase in spontaneous action potential frequency in *App*$^{NL\text{-}G\text{-}F}$ locus coeruleus neurons. The nucleus loses cells and raises output.

Firing is what removes the axon. This is the causal step established in Section III and it bears repeating in its strongest form, because everything in this section follows from it. In Meyer's account the sequence is: increased spontaneous firing → calcium entry → calcium-dependent externalisation of phosphatidylserine on the axonal membrane → MFG-E8 bridging → microglial recognition and phagocytosis of the axon. Blocking the phagocytic step preserved the axon and the function. The axon is removed as a consequence of the activity of the neuron that owns it.

If that is right, then the second trigger is a control failure, and one should ask what normally controls the firing rate of an autonomously active neuron. The answer, in every autonomous pacemaker that has been studied, is potassium.

KCNIP4 and the A-type brake. The Kv4 family of voltage-gated potassium channels carries the A-type current — a low-threshold, rapidly activating and rapidly inactivating outward current that opposes depolarisation in the subthreshold range, and which therefore sets the interval between spikes in a pacemaking neuron. Kv4 channels do not operate alone: their kinetics are set by auxiliary subunits, of which the KChIP family, encoded by the *KCNIP* genes, is the principal one. *KCNIP4* encodes KChIP4, first cloned as CALP by Morohashi and colleagues (2002) — who identified it, as it happens, through its interaction with the carboxy-terminal region of presenilin 2, so that the gene enters the Alzheimer literature by way of a familial Alzheimer protein rather than by way of excitability at all.

The functional point is that KChIP proteins are not merely accessories but determinants of how much brake the neuron has. Increased binding of KChIP4 enhances the recovery of Kv4.2 from inactivation, and the KChIP4a splice variant, which carries a thirty-four-residue K-channel-inactivation-suppressor domain at its amino terminus, dramatically slows Kv4 inactivation — extending time to half-inactivation from roughly 30 milliseconds for Kv4.2 alone to roughly 100 milliseconds. A brake that stays engaged three times longer is a substantially different brake. And the isoform that supplies it differs from its siblings almost entirely in a domain generated by alternative splicing — a fact Section VII will make use of.

The human evidence that KCNIP4 is a resilience factor. Until recently the argument would have stopped at plausibility. It no longer does. Dharshini and colleagues (2026), using single-nucleus and spatial transcriptomics to compare neocortical regions affected early in Alzheimer's disease (prefrontal cortex, precuneus) with one affected late (primary visual cortex), identified a resilient excitatory population in layer 4 of the primary visual cortex and asked what distinguishes it. *KCNIP4* was consistently upregulated in the early stages of pathology in the resilient neurons, alongside a programme of synapse-maintenance and calcium-homeostasis genes; it was downregulated in vulnerable neurons during the stages of cell death, and declined overall in late disease. The authors then did the experiment that converts a correlation into a mechanism: adeno-associated-viral overexpression of *Kcnip4* in excitatory cortical neurons reduced the activity-dependent genes *Arc* and *c-Fos*, which they interpret as a compensatory mechanism against neuronal hyperexcitability.

Read against the coerulean chronology, this is a striking convergence. A gene whose product damps neuronal activity is upregulated *early* in the neurons that survive, and fails in the neurons that die. If the coerulean axon is removed because the coerulean neuron fires too much, then the presence or absence of this brake is a candidate answer to why some coerulean neurons lose their arbor and their lives on schedule and others do not.

The proof of principle from next door. The strongest support for the general shape of this argument comes not from the locus coeruleus but from a neighbouring wake-promoting nucleus,

and the honest thing is to present it as an analogy rather than as evidence about the locus coeruleus. Li and colleagues (2022) showed that hypocretin/orexin neurons of the lateral hypothalamus become hyperexcitable in aged mice, that the hyperexcitability is attributable to lower *KCNQ2* expression and an impaired M-current, that this drives fragmentation of sleep, and — decisively — that disrupting *Kcnq2/3* in the hypocretin neurons of young mice reproduces the aged phenotype. An arousal-promoting neuron loses a potassium brake; it fires more; sleep fragments.

The channel is different (KCNQ, not Kv4), the auxiliary subunit is not the same, and the nucleus is not the locus coeruleus. What transfers is the architecture: **in a wake-promoting nucleus, a potassium brake is the thing that fails with age, and its failure is sufficient to produce the phenotype.** That is a precedent, not a proof, and Section XIII grades it as such.

The proposal, stated plainly. The second trigger that Theofilas's data demand is, on this reading, the loss of excitability control in a neuron that was already sick with tau — and its first casualty is not the cell body but the arbor, which is removed by microglia reading a calcium-driven eat-me signal that the neuron's own firing generates. Tau makes the neuron sick and keeps it sick for twenty-five years. Hyperexcitability makes the sickness lethal to the axon. Only much later, and by mechanisms this dissertation does not claim to have identified, does the denuded soma die.

VII. The Jammed Nucleus — Tau, Pinin, and the Splicing of the Brake

A gap remains in the account just given. If the tau clock and the axon clock are separate, what connects them? Why should a neuron that has carried pretangle tau for two decades begin, at some point, to lose control of its firing? This section proposes a candidate connection, states clearly that it is an inference rather than a measurement, and identifies the experiment that would test it.

A note on an unfortunate abbreviation. The protein at the centre of this section is pinin, the product of the human gene *PNN*. In the literature of this corpus, "PNN" almost invariably denotes the perineuronal net, an extracellular lattice of lecticans and hyaluronan that has been the subject of three companion volumes. These are entirely different objects: one is a sulfated extracellular matrix; the other is an intranuclear protein of the splicing machinery. They share three letters and nothing else. To avoid a collision that would be actively misleading, this dissertation writes **pinin** throughout for the protein and reserves "perineuronal net" for the matrix. That the locus coeruleus lacks the one and loses the other is a coincidence of nomenclature, not a mechanism, and no argument here depends on the pun.

Tau disassembles the nuclear speckle. Nuclear speckles are membraneless organelles that concentrate the machinery of pre-messenger-RNA splicing. Lester and colleagues (2021) established that tau aggregates are not merely proteinaceous but are RNA–protein assemblies, enriched for small nuclear and small nucleolar RNAs, and that nuclear tau aggregates colocalise with nuclear speckles and alter their composition, dynamics and spatial organisation. Several speckle components mislocalise to cytosolic tau aggregates — in cells, in mouse brain, and in the brains of individuals with Alzheimer's disease, frontotemporal dementia and corticobasal degeneration. The scaffold protein SRRM2 is the most prominent of these, and it drags its partners with it: the authors name **pinin**, along with SFPQ and DYRK1A, as proteins depleted from nuclear speckles by association with cytosolic tau aggregates.

The functional consequence is large and was measured. In cells bearing tau aggregates, splicing failed extensively: the retention of more than 1,200 introns encoded by 641 genes, correlating with the altered speckle properties. This is not a subtle shift in isoform ratios. It is the partial disassembly of the apparatus by which a neuron decides which version of each of its proteins to build.

Why this belongs in the coerulean interval. Place this result on the chronology of Section V. Tau is present in coerulean neurons from Braak stage 0 and in an increasing fraction thereafter, while those neurons remain alive and countable for another twenty-five years. The question of what

tau is *doing* to them during that time has had no good answer; "sickening them" is a description, not a mechanism. Lester's result supplies a mechanism of exactly the right kind — one that degrades the neuron's function progressively and profoundly without killing it, and that scales with tau burden. A neuron whose speckles have been stripped of SRRM2 and pinin is a neuron that is still alive, still countable, still tyrosine-hydroxylase-positive, and no longer able to splice reliably. That is a precise description of the occupant of the interval.

The inferential bridge, and its status. Now the speculative step, which is stated as such. *KCNIP4* is a gene whose functional output depends on alternative splicing to an unusual degree: the several KChIP4 variants differ from one another principally in the amino-terminal domain, and it is precisely that domain which distinguishes KChIP4a — the isoform bearing the inactivation-suppressor sequence that triples the persistence of the A-type brake — from variants that lack it. A neuron that cannot splice reliably is a neuron whose ratio of braking to non-braking KChIP4 isoforms is no longer under control.

The inference, therefore, is this: **tau, by disassembling the nuclear speckle and depleting pinin and SRRM2 from it, degrades the splicing fidelity on which the neuron's excitability brake depends; the brake weakens; the pacemaker accelerates; and the accelerated pacemaker's calcium load flags its own axon for phagocytic removal.** This would make the tau clock and the axon clock not merely consecutive but causally linked, with the splicing lesion as the transmission between them — and it would identify the second trigger as the point at which splicing degradation crosses the threshold at which excitability control fails.

The status of this proposal must be stated without softening. Each link in it is individually evidenced — tau depletes pinin from speckles and causes mass intron retention (measured, Lester 2021); KChIP4 isoforms are splice-determined and differ in braking power (measured, in heterologous expression); *KCNIP4* is upregulated in resilient neurons and its overexpression damps activity markers (measured, Dharshini 2026); coerulean neurons fire more in the diseased state and that firing removes their axons (measured, Meyer 2025). But **no one has measured** *KCNIP4* **splicing in locus coeruleus neurons in Alzheimer's disease**, and until someone does, the chain is a hypothesis with four sound links and one unmade joint. It is offered here because it is testable, and Section XIV names the test.

One caution against over-reading. Intron retention across 641 genes is a shotgun, not a rifle. If splicing failure degrades the excitability brake it will also degrade a great deal else, and it would be a mistake to present *KCNIP4* as the uniquely important casualty rather than as the casualty most relevant to the argument of this volume. The honest formulation is that the splicing lesion is a plausible general mechanism of progressive coerulean dysfunction, and that excitability control is one specific function it would be expected to compromise.

VIII. The Executioners — C4d, LilrB2, and the Tag on the Synapse

Section III established that the coerulean axon is removed by microglia reading phosphatidylserine, and recorded the direct finding that complement C1q is *not* differentially deposited on those axons. Complement nonetheless has a well-evidenced place in this chronology. The place is the synapse, and the relevant fragment is C4d.

What C4d is. When the classical complement pathway is activated, C4 is cleaved and a fragment becomes covalently attached to nearby surfaces. C4d is the terminal split product of that process, and its defining property is durability: because it is covalently bound and not readily removed, C4d is used in transplant pathology precisely as a persistent record that complement was activated at a site. It is, in the strictest sense available in immunopathology, a *timestamp*. In Alzheimer's disease brain, activated fragments C3d and C4d are strongly detected, and C4d immunostaining marks senile plaques, diffuse amyloid deposits, dystrophic neurites and some neurofibrillary tangles.

What C4d does. Until recently C4d was regarded as a footprint — evidence that something had happened, not an agent that made something happen. Brott and colleagues (2025) changed this. They showed that C4d binds LilrB2, the human receptor whose murine homolog is PirB, with nanomolar affinity; that C4d and LilrB2 colocalise at excitatory synapses in human cerebral cortex by array tomography, and colocalise with amyloid-beta in Alzheimer's disease; that both C4 and C4d increase with age and increase further in Alzheimer's disease; and that C4d mediates pruning and elimination of dendritic spines on cortical pyramidal neurons. C4d is therefore not only a record of complement activation but a high-affinity ligand that instructs synapse removal through a specific receptor.

The receptor matters for this corpus. LilrB2/PirB is the receptor through which amyloid-beta oligomers were shown to impair synaptic plasticity in Alzheimer models (Kim and colleagues, 2013), and the complement pathway's role in early synapse loss in those models is established (Hong and colleagues, 2016), as is the developmental precedent that made the field take complement seriously in the first place — the schizophrenia risk conferred by structural variation at the *C4* locus, acting through synaptic pruning (Sekar and colleagues, 2016). Brott's result unifies these: one receptor, two ligands, both elevated in the disease, both instructing the same removal.

Where it sits on the coerulean clock. The coerulean neuron's synapses are at the far end of the arbor whose loss constitutes clock two, and the natural reading is that synaptic removal and

axonal removal are two phases of the same dying-back. But the tags are demonstrably different — phosphatidylserine and MFG-E8 on the axon, C4d and LilrB2 at the synapse — and it would be a mistake to collapse them. The more defensible statement is that the coerulean projection is disassembled from its distal end inward by at least two independent, microglially-executed, non-cell-autonomous mechanisms, each with its own molecular tag, and that a neuron losing terminals to one and axon segments to the other has no route by which its cell body can defend either.

There is also a plainer reason to place C4d here rather than at the coerulean axon, and it is a reason about norepinephrine. Norepinephrine is a suppressor of microglial inflammatory activation: Heneka and colleagues (2010) showed that coerulean lesion increases amyloid pathology by de-repressing microglial function, and the same axis is the subject of the companion volume *The Corruption of the Gardener*. A locus coeruleus that is losing its projection is a locus coeruleus that is withdrawing the noradrenergic restraint from cortical microglia. Those disinhibited microglia are the cells that read C4d and prune. **The coerulean axon's removal therefore lifts the brake on the very cells that remove cortical synapses** — which is the first of two feedback loops this dissertation identifies, and the reason clock two, once running, is difficult to stop.

IX. The Ratchet — Ephexin-5 and the Synapse That Cannot Return

A synapse lost is not necessarily a synapse gone. Under normal conditions the excitatory synapse is a renewable structure, formed and eliminated continuously, and a loss that is not accompanied by a block on replacement is a loss that a healthy neuron can make good. This section places the molecule that blocks replacement.

Ephexin-5 is a developmental brake on synapse formation. Margolis and colleagues (2010) identified Ephexin-5 as a guanine nucleotide exchange factor for the small GTPase RhoA that negatively regulates excitatory synapse development — a restraint that holds synapse number down until it is lifted. The lifting is done by EphB receptor signalling: ephrin-B binding to the EphB2 receptor tyrosine kinase triggers phosphorylation of Ephexin-5, its ubiquitination, and its proteasomal degradation. The ubiquitin ligase responsible is Ube3A, mutated in Angelman syndrome — which is how the molecule first acquired clinical significance and why its logic is unusually well characterised. Ephexin-5, in short, is a brake whose normal fate is destruction: the synapse forms when Ephexin-5 is removed.

In Alzheimer's disease it is not removed. Sell, Schaffer and Margolis (2017) showed that amyloid-beta acutely promotes Ephexin-5 production in mature hippocampal neurons and in mice expressing human amyloid precursor protein, and that Ephexin-5 expression is highly elevated in the hippocampi of human Alzheimer's disease patients. Genetic removal of Ephexin-5 prevented the hippocampal spine-density abnormalities and the cognitive deficits in the *hAPP* mice. Two arms of the same lesion converge: amyloid raises Ephexin-5 directly, and EphB2 — the kinase whose signalling would degrade it — is itself depleted in Alzheimer's disease, its restoration sufficient to rescue cognitive function in a mouse model (Cissé and colleagues, 2011). The brake is pushed on from one side and its release is disabled from the other.

Why this is the ratchet. The functional significance of Ephexin-5 for the present argument is not that it destroys synapses. It does not; C4d and complement-directed microglial pruning do that, and amyloid does it, and the withdrawal of noradrenergic support does it. What Ephexin-5 does is make the destruction permanent. A neuron that has lost spines and has an elevated RhoA exchange factor restraining excitatory synapse formation is a neuron that cannot rebuild what it lost. The loss is converted from a deficit into a floor.

This has a specific consequence for the coerulean chronology. Section III showed that the surviving coerulean neurons *sprout* — Szot and colleagues documented dendritic sprouting into the

peri-coerulean zone and axonal sprouting to the hippocampus in Alzheimer and Lewy-body tissue, and Kelly's transcriptional data showed the machinery of neuritic structural plasticity being downregulated as the disease advances. The system tries to rebuild, and then stops being able to. Ephexin-5 is one named molecular reason the rebuilding stops, operating in the target tissue rather than in the locus coeruleus itself, and it converts the axon clock from a process that might in principle reverse into one that only ratchets in one direction.

A limitation should be recorded here rather than deferred. All of the Ephexin-5 evidence in Alzheimer's disease concerns hippocampal excitatory synapses and amyloid-driven models. No one has measured Ephexin-5 at noradrenergic terminals, in the locus coeruleus, or in relation to coerulean denervation. Its inclusion in this chronology is a statement about the *target field* into which the coerulean axon projects and in which its terminals are lost — an environment that is hostile to synaptic rebuilding — and not a claim about the coerulean neuron's own biology. Section XIII grades it accordingly.

X. REM and NREM — The Interval in Which the Clock Slows

If the coerulean axon is destroyed by the coerulean neuron's own firing, then any interval during which that neuron does not fire is an interval during which the axon is not being destroyed. The brain contains exactly one such interval, it recurs several times a night, and it is called REM sleep. This section argues that sleep is not a consequence of coerulean pathology to be catalogued alongside the others, but the tempo control on clock two — and that the locus coeruleus is asked to do two opposite things in the two sleep states, so that a sickening nucleus fails in two distinct ways.

The locus coeruleus is the brain's only neuron population with a scheduled total off-period. Aston-Jones and Bloom established in 1981 that coerulean discharge is state-dependent across the sleep-wake cycle and, uniquely, that it ceases in REM. The canonical rates from that work and its successors are approximately 2.12 Hz in waking, 0.69 Hz in slow-wave sleep, and 0.02 Hz in REM — which is to say, effectively zero. Coerulean neurons are "REM-off" cells in the strict sense. Nothing else in the arousal system is granted an off-switch this complete.

Set this beside Meyer's mechanism. Firing drives calcium entry; calcium entry drives phosphatidylserine externalisation on the axolemma; externalised phosphatidylserine is read by microglia as an instruction to remove the process. Then REM sleep is the only period in the twenty-four-hour cycle in which that instruction is not being written. **REM is the coerulean axon's reprieve** — the interval in which calcium load falls to floor, the eat-me flag is not being generated, and whatever membrane repair the neuron is capable of can proceed against a quiet background. On this reading the fragmentation and loss of REM sleep that accompanies ageing and Alzheimer's disease is not a symptom of coerulean damage downstream of it. It is a withdrawal of the axon's protection, and therefore a lesion in its own right.

NREM asks the opposite. The complementary state makes a demand that is not silence but rhythm. Hauglund and colleagues (2025) identified tightly synchronised oscillations of norepinephrine, cerebral blood volume and cerebrospinal fluid as the strongest predictors of glymphatic clearance during NREM sleep. Optogenetic stimulation of the locus coeruleus induced anti-correlated changes in vasomotion and cerebrospinal-fluid signal; driving arterial oscillations enhanced cerebrospinal-fluid inflow; vasomotion, in their account, is the pump. And the pharmacological control is the most telling result of the study: zolpidem, a widely prescribed hypnotic, suppressed the norepinephrine oscillations and suppressed glymphatic flow with them. Sleep that lacks the coerulean oscillation does not clear the brain, even though it is, by the

electroencephalogram, sleep.

So the locus coeruleus is required to *oscillate slowly* through NREM and to *fall silent* through REM. It is a two-state obligation, and a nucleus in which pacemaking has become deranged upward — Section VI — fails both halves at once. It cannot generate a clean infra-slow oscillation because its baseline firing is elevated and irregular; it cannot deliver a true REM silence because the population that must fall silent is fragmenting and the sleep architecture that schedules the silence is itself breaking down.

The human evidence, and it is early-stage evidence. Van Egroo and colleagues (2021) imaged the locus coeruleus at 7 Tesla in seventy-two cognitively unimpaired older individuals aged fifty to eighty-five and found that worse coerulean structural integrity was associated with more frequent nocturnal awakenings, in the context of plasma markers of neurodegeneration. The subjects were *cognitively unimpaired*. This is the interval — the decades between the tau and the death — and the coerulean-sleep coupling is already measurable in it. Oh and colleagues' finding that wake-promoting nuclei sustain substantial neuronal loss specifically in Alzheimer's disease supplies the neuropathological counterpart, and their recommendation is worth taking literally: degeneration of these nuclei "should be included in the models explaining sleep-wake disturbances in AD" — that is, the sleep disturbance of Alzheimer's disease is in significant part a *brainstem lesion*, not merely a behavioural or circadian one.

The second feedback loop. Section VIII identified the first: coerulean axon loss withdraws noradrenergic restraint from microglia, and disinhibited microglia remove more of the projection. The sleep axis supplies the second, and it is tighter. Coerulean pathology fragments sleep. Fragmented sleep removes the REM interval during which the coerulean axon is not being flagged for phagocytosis, and degrades the NREM oscillation that clears the extracellular space of the amyloid and tau species that the companion volumes have traced. More tau and less reprieve produce more coerulean pathology. The nucleus that governs the restorative state is destroyed by the failure of the state it governs.

This is the third instance in this corpus of what the companion volumes have called the arsonist operating the sprinkler — the locus coeruleus as both the source of the damage and the operator of the system that would repair it — and it is the instance with the clearest therapeutic implication, because sleep is modifiable and pacemaker excitability is druggable.

A caution about the obvious intervention. It does not follow from any of this that sedation is protective. Hauglund's zolpidem result is a direct demonstration that a drug can produce sleep by the electroencephalogram while abolishing the coerulean oscillation that makes sleep useful. A hypnotic that silences the locus coeruleus would, on the argument of this section, protect the axon and abolish the clearance — buying the REM benefit at the cost of the NREM one. The intervention the model actually nominates is not suppression but *restoration of the alternation*: an intact

REM silence and an intact NREM oscillation, which is a considerably harder pharmacological brief and one that current hypnotics do not meet.

XI. Reelin — The Brake That Was Never Staged, and the Next Domino

Reelin enters the coerulean chronology twice, at opposite ends, and neither entry is the one the reelin literature would predict.

First entry: the brake the locus coeruleus never had. The companion volumes of the reelin series established the following chain. Reelin, acting through ApoER2 and VLDLR, induces tyrosine phosphorylation of Disabled-1, which recruits phosphoinositide-3-kinase and activates Akt, which phosphorylates glycogen-synthase-kinase-3-beta on its regulatory serine and holds the principal tau kinase in its inhibited state. The signal is a continuous inhibitory input to the tau kinase — a brake. And the brake cannot be applied in free solution: it requires N-sulfated heparan sulfate as an obligate co-receptor, presented on a sulfated extracellular surface, and the condensed form of that surface is the perineuronal net.

Morawski and colleagues established that the subcortical nuclei attacked earliest and hardest by tau — the locus coeruleus foremost, with the nucleus basalis, the raphe and the dorsal tegmentum — are precisely those devoid of an aggrecan-based perineuronal net, while net-ensheathed neurons in the thick of the pathology rarely tangle. Matchett and colleagues list the absence of the net among the mechanisms of coerulean selective vulnerability.

The consequence for this dissertation's chronology is a subtraction rather than an event. In the cortex, the reelin brake is *cut* — by the proteolytic and inflammatory mechanisms traced in *The Coerulean Shears* and *The Speck and the Architect*, which strip the sulfated bed and shed the receptors. In the locus coeruleus there is nothing to cut. The nucleus begins life without the staging surface on which the brake would be presented, which is why it needs no second event to be vulnerable and why its tau clock starts in the third decade rather than in the seventh. **The locus coeruleus is not a structure that loses its protection during the disease. It is a structure that enters the disease already without it.**

This is worth stating because it is discouraging in a specific way. Interventions aimed at preserving or restoring the sulfated matrix — the therapeutic direction that the reelin and perineuronal-net volumes converge on — would, if they worked perfectly, do nothing for the locus coeruleus, because the locus coeruleus has no such matrix to preserve. The nucleus that fails first is the nucleus least reachable by the corpus's most-developed protective strategy. Any account that claims otherwise is claiming more than the anatomy allows.

Second entry: reelin marks the next domino. The other end of the reelin story places it not in the locus coeruleus but at the structure that receives the coerulean lesion next. Kobro-Flatmoen and colleagues (2016) showed that in layer II of the entorhinal cortex — the cortical population that fails earliest in Alzheimer's disease, and the target of the earliest coerulean projections — the neurons that selectively express intracellular amyloid-beta in early disease are the **reelin-immunoreactive** ones. The association is present in a transgenic rat model at the pre-plaque stage and in human subjects with early Alzheimer-related pathological change. Subsequent work from the same group showed that lowering reelin levels in these neurons lowers their intracellular amyloid-beta.

So reelin identifies, by immunoreactivity, the specific cortical population that stands second in the anatomical sequence after the locus coeruleus. This is a marker relationship and not, on present evidence, a causal one in the direction that would be most convenient; indeed the finding that lowering reelin lowers intracellular amyloid-beta cuts against a simple protective reading and is one of the reasons the reelin literature contains a genuine and unresolved tension. What the finding establishes for the present purpose is narrower and firmer: the two structures at the head of the disease's anatomical sequence — the locus coeruleus and entorhinal layer II — are linked by reelin in opposite senses. The first lacks the surface on which reelin's brake would be staged. The second is defined, at the level of which cells go first, by expressing reelin itself.

What reelin does not explain. It should be said plainly that reelin does not explain the timing this dissertation is about. The absence of a perineuronal net in the locus coeruleus is a constitutional fact, present from birth and unchanged through life. A constant cannot explain a change. It can explain why the coerulean neuron is the first to accumulate tau; it cannot explain why that neuron, having carried tau for twenty-five years, begins at some particular point to lose its axon and later its life. The second trigger must be something that changes, and reelin's absence does not change. This is the reason Sections VI and VII look to excitability and splicing rather than to the matrix — and it is a case where following the corpus's dominant thread would have produced a worse answer than departing from it.

XII. How Coerulean Decline Accelerates Death Elsewhere

The three clocks describe what happens to the locus coeruleus. This section addresses the reciprocal question named in this volume's brief: what the coerulean decline does to the rest of the brain, and specifically how it accelerates neuronal death beyond the nucleus.

The controlled experiment. The cleanest causal evidence comes from lesion studies in tau transgenic animals. Chalermpalanupap and colleagues (2018) ablated the locus coeruleus with the noradrenergic neurotoxin DSP-4 in P301S tau transgenic mice beginning at two months of age. The lesioned animals showed significantly impaired hippocampal-dependent contextual fear memory — a deficit detectable *only* in the animals that had both tau pathology and coerulean ablation, indicating synergy rather than addition — and, decisively for this section, accelerated cell death and increased mortality.

The result that gives this experiment its analytic force is a negative one: the acceleration occurred **without profoundly affecting forebrain tau pathology**. Removing the locus coeruleus made neurons die faster without making more tangles. Whatever the coerulean projection is contributing to the survival of its targets, it is not contributing by restraining tau, and the standard framing in which the locus coeruleus matters because it modulates proteinopathy is, on this evidence, incomplete. It matters because it keeps cells alive.

The microglial route. The best-characterised mechanism for that trophic role is immunological. Heneka and colleagues (2010) showed that the locus coeruleus controls Alzheimer's disease pathology by modulating microglial function through norepinephrine: coerulean lesion in an amyloid model increased pathology by removing noradrenergic suppression of microglial inflammatory activation. Norepinephrine is, in this account, a standing restraint on the brain's resident immune cell — and the companion volume *The Corruption of the Gardener* grades the noradrenergic axis among the threshold-setting drivers of microglial dysfunction, distinguished not by being the largest such driver but by being the earliest and the most pharmacologically reachable.

Combine this with Section VIII and the arithmetic of clock two becomes unpleasant. The coerulean axon is what delivers norepinephrine to cortical microglia. Clock two removes the axon at Braak stage I–II. Therefore the noradrenergic restraint on cortical microglia is withdrawn at Braak stage I–II — decades before any coerulean cell body dies, and long before any imaging measure of nuclear integrity would report a catastrophe. The disinhibited microglia then prune C4d-tagged synapses and phagocytose phosphatidylserine-tagged axons, including further coerulean ones. The

loop closes.

The clearance route. Section X supplies the second mechanism. If the infra-slow noradrenergic oscillation of NREM sleep is the pump that drives glymphatic clearance, then a degraded coerulean output degrades clearance, and a brain that clears less accumulates more of the species that damage it. This is a slower loop than the microglial one and its evidence base in humans is thinner — the 2024–2025 controversy over whether clearance genuinely increases in sleep has not been settled, and the companion volume *The Glymphatic Collapse* treats it at length — but its direction is not seriously in doubt.

The reserve route. The third mechanism is the least mechanistic and the most clinically grounded. Wilson and colleagues (2013) found coerulean noradrenergic neuronal density to be associated with the rate of cognitive decline across nearly six years of annual assessment in 165 autopsied participants, and framed the nucleus as a structural component of neural reserve. Two people with the same cortical pathology decline at different rates, and one of the things that differs between them is how much locus coeruleus they have left. This is not a mechanism, but it is the outcome any proposed mechanism has to produce.

Why the acceleration is not simply "less norepinephrine." It is worth resisting the deficiency reading one last time. The interval in which the coerulean projection is being lost is also the interval in which cerebrospinal-fluid norepinephrine is *rising* with disease severity (Elrod and colleagues, 1997), because the surviving neurons compensate — raising synthesis, sprouting, and firing harder (Szot and colleagues, 2006). The target tissue is therefore not simply deprived; it is deprived of *regulated, spatially precise, state-dependent* norepinephrine and simultaneously exposed to a raised and deranged tonic level. The companion volume *The Coerulean Pincer* argues that this excess is itself injurious, driving matrix proteolysis from outside the target neuron and over-activating glycogen-synthase-kinase-3-beta within it. The present volume adds the observation that the same excess is what removes the axon delivering it. Compensation is not free, and its bill is paid in arbor.

XIII. The Validity Ledger

Every claim load-bearing in this dissertation is graded below. Tier I denotes direct measurement in human tissue or a well-controlled human cohort. Tier II denotes direct measurement in an animal model or *in vitro* system with a clear human correlate. Tier III denotes inference assembled across literatures that do not cite one another — the claims that are original to this volume and that carry the corresponding risk.

Claim	Tier	Basis	What would overturn it
Tau appears in the locus coeruleus in the third decade, before anywhere else	I	Braak & Del Tredici 2011; Braak et al. 2011 (2,332 brains)	Detection of an earlier lesion elsewhere with equivalent method
7.9% of LC neurons bear hyperphosphorylated tau inclusions at Braak stage 0	I	Ehrenberg et al. 2017, unbiased stereology, n=48	Replication failure in an independent stereological series
LC neuronal number is preserved through Braak 0–II and falls from III	I	Theofilas et al. 2017, unbiased stereology across full Braak range	A series with adequate power finding loss at stage I–II
~30% LC neuron loss is already present at amnestic MCI	I	Kelly et al. 2017, TH stereology	Replication failure; or evidence aMCI ≠ Braak III–IV in these cohorts
End-stage LC loss is 63–83%	I	Zarow et al. 2003 (n=86 AD); Matchett et al. 2021 review	— (not in dispute)
LC neuronal loss is AD-specific among tauopathies	I	Oh et al. 2019 (AD vs. PSP vs. CBD)	Larger PSP/CBD series showing equivalent loss
Noradrenergic fibre loss in a target field is present at Braak 1–2 and does not worsen after	I (human OB)	Meyer et al. 2025, human olfactory bulb, NET immunostaining	Failure to replicate in an independent bulb series
Axon loss precedes and dissociates from LC cell-body loss	II	Meyer et al. 2025: 33% fibre loss at 6 mo, no neuron loss at 12 mo, *App*$^{NL\text{-}G\text{-}F}$	Demonstration that mouse fibre loss reflects transport/marker failure, not axon loss
LC axon removal is microglial, via phosphatidylserine and MFG-E8, not complement	II	Meyer et al. 2025 (C1q colocalisation unchanged)	Evidence of complement-dependent LC axon removal in human tissue

Claim	Tier	Basis	What would overturn it
LC neurons fire *more* in the diseased state; firing drives axonal PS externalisation	II	Meyer et al. 2025 electrophysiology; Ca2+-dependence	Recordings showing LC hypoactivity in early human disease
Surviving human LC neurons compensate (TH mRNA up, sprouting)	I	Szot et al. 2006, post-mortem AD and DLB	—
CSF norepinephrine rises with AD severity	I	Elrod et al. 1997	Larger contemporary CSF cohorts reversing the direction
LC neuronal density is a component of neural reserve	I	Wilson et al. 2013, n=165, 5.8 yr follow-up to autopsy	—
C4d binds LilrB2 with nanomolar affinity and mediates spine pruning; both rise in AD	I/II	Brott et al. 2025, human cortex array tomography + functional assay	Failure to replicate the pruning phenotype
Ephexin-5 is elevated in human AD hippocampus; its removal rescues *hAPP* mice	I/II	Sell et al. 2017; Margolis et al. 2010	—
Tau aggregates deplete pinin and SRRM2 from nuclear speckles; >1,200 introns retained	I/II	Lester et al. 2021, cells + mice + human AD/FTD/CBD brain	—
KCNIP4 is upregulated early in resilient neurons; overexpression damps Arc/c-Fos	I/II	Dharshini et al. 2026, snRNA-seq + spatial + AAV	Failure to replicate the resilience association
LC neurons are REM-off; ~2.1 Hz wake, ~0.7 Hz SWS, ~0.02 Hz REM	II	Aston-Jones & Bloom 1981 and successors, rodent	—
NREM noradrenergic oscillation drives glymphatic clearance; zolpidem abolishes both	II	Hauglund et al. 2025, mouse, optogenetic + pharmacological	The unresolved 2024–25 sleep-clearance controversy resolving against it
Worse LC integrity tracks nocturnal awakenings in unimpaired older adults	I	Van Egroo et al. 2021, 7T MRI, n=72	—
LC ablation accelerates neuron death and mortality without increasing forebrain tau	II	Chalermpalanupap et al. 2018, DSP-4 in P301S	—

Claim	Tier	Basis	What would overturn it
LC lacks a perineuronal net; net-bearing neurons resist tangles	I	Morawski et al. 2010; Matchett et al. 2021	—
Reelin-immunoreactive EC layer II neurons selectively carry early intracellular Aβ	I	Kobro-Flatmoen et al. 2016, rat model + human	—
The "second trigger" is loss of excitability control	**III**	**Inference: Theofilas's gap + Meyer's firing-driven axon loss + Dharshini's KCNIP4 resilience**	**Demonstrating LC axon loss is firing-independent**
Tau→speckle disassembly→splicing failure degrades the KCNIP4-dependent brake	**III**	**Inference across Lester 2021 + KChIP4a splice biology + Dharshini 2026. Untested joint.**	**Measuring KCNIP4 splicing in AD locus coeruleus and finding it unchanged**
The whole LC arbor follows the olfactory bulb's schedule	**III**	**Extrapolation from one target field, supported by the rostrocaudal gradient**	**Cortical/hippocampal NET fibre quantification at Braak I–II showing no loss**
Ephexin-5 acts as a ratchet on coerulean target fields specifically	**III**	**Extrapolation: all Ephexin-5 AD data are hippocampal/amyloid-driven; none noradrenergic**	**Measuring Ephexin-5 at noradrenergic terminals and finding it unchanged**
REM silence is protective for the coerulean axon	**III**	**Inference: Aston-Jones REM-off + Meyer firing-dependence. Never tested jointly.**	**REM deprivation failing to accelerate LC axon loss in a model**

The four Tier III rows are this dissertation's original contribution and its exposure. They are stated as hypotheses with named tests, not as findings.

XIV. What Would Falsify the Three-Clock Model

A model that cannot be killed is not worth stating. The three-clock account makes four predictions that are wrong if the model is wrong, and each is measurable with existing methods.

One. Cortical noradrenergic fibre density should already be substantially reduced in human Braak stage I–II tissue. The strongest human evidence for clock two comes from a single target field, the olfactory bulb. If the model is right, norepinephrine-transporter immunostaining in prefrontal and entorhinal cortex from Braak I–II donors should show fibre loss of comparable magnitude — on the order of a quarter to a third — in the presence of normal coerulean cell counts in the same brains. If cortical fibre density at Braak I–II is normal, then the olfactory bulb is a special case, clock two is not general, and the axon-first claim collapses to a claim about olfaction.

Two. *KCNIP4* splicing should be disrupted in Alzheimer locus coeruleus neurons, and disrupted before cell loss. This is the untested joint in Section VII, and it is directly addressable: laser-capture or single-nucleus isolation of coerulean neurons across Braak stages, with isoform-resolved sequencing, asking whether the KChIP4a-type inactivation-suppressor-containing isoform declines relative to its siblings, and whether the decline precedes the stage-III inflection in cell number. If *KCNIP4* isoform ratios are stable across the interval, the splicing bridge is dead and the second trigger must be sought elsewhere.

Three. Suppressing coerulean firing should preserve the coerulean axon. The prediction that most sharply distinguishes this model from the deficiency model is that *quieting* the locus coeruleus should protect its projection. Chemogenetic or pharmacological reduction of coerulean pacemaker rate in the *App*$^{NL\text{-}G\text{-}F}$ mouse should reduce phosphatidylserine externalisation and preserve norepinephrine-transporter fibre density at six months. The deficiency model predicts the opposite — that quieting the nucleus should worsen everything. This is a clean, cheap, decisive experiment, and to the best of the present survey it has not been done.

It should be said that the same experiment carries a real risk of a split result, and that the split would be informative rather than disappointing: axon preserved, cognition worsened. Rorabaugh and colleagues (2017) restored reversal learning in the TgF344-AD rat by chemogenetic *activation* of the locus coeruleus, and a good deal of the noradrenergic therapeutic literature points the same way. If suppression preserves the arbor while impairing the cognition the arbor exists to support, then the model is right about the mechanism and the therapeutic implication is a timing problem

rather than a direction — which is precisely the sort of finding this corpus exists to surface.

Four. REM deprivation should accelerate coerulean axon loss. If REM silence is the axon's reprieve, selective REM deprivation in an amyloid model should accelerate the fibre loss that Meyer and colleagues timed, and REM enhancement should retard it, with both effects tracking calcium load and phosphatidylserine exposure rather than tau burden. A null result would sever sleep from the axon clock and reduce Section X to an association.

There is a fifth prediction that is not falsifiable with present technology and is recorded as an aspiration rather than a test: that in a living human, coerulean *axonal* integrity and coerulean *somatic* integrity should be separable measurements, and that the first should decline decades before the second. Diffusion-based measures of the coerulean tract are beginning to approach this, but the resolution is not yet adequate to make the dissociation this dissertation asserts directly visible in a living person.

XV. Conclusion — Where in the Interval to Intervene

The question this dissertation was set was one of comparative timing: when do locus coeruleus cells die, when do their axons degenerate, and how do these relate to the chronology of the disease. The answer is that they are three different questions with three different answers separated by decades, and that the field's habit of treating them as one has directed attention to the least useful of the three.

The cell bodies die late. Not at the beginning of the disease, where the tau is, but from Braak stage III — by which point, translated into clinic terms, roughly a third of the nucleus is already gone and the patient is presenting with amnestic mild cognitive impairment. The somatic clock is therefore not an early-detection instrument and not a therapeutic target; by the time it is running, the process it would be asked to interrupt has been under way for twenty-five years.

The axons degenerate early. A third of the projection to a coerulean target field is lost in human tissue at Braak stage 1–2, and in the corresponding mouse the fibres are a third gone at six months while the cell bodies are still uncounted-for-loss at twelve. This is the lesion that matters, because it is the lesion that removes the function — the noradrenergic restraint on microglia, the state-dependent modulation of the cortex, the oscillation that drives clearance — while every measurement the field routinely makes still reports an intact nucleus.

And the interval between them is occupied. It is not a latency during which nothing happens. Inside it, tau disassembles the nuclear speckle and depletes pinin, and splicing fidelity degrades across hundreds of genes; the neuron's excitability brake, which is isoform-dependent and which the resilient neurons of the cortex are now known to upregulate, plausibly fails; the pacemaker accelerates; calcium floods the arbor; phosphatidylserine appears on the axolemma; microglia — themselves progressively released from the noradrenergic restraint that the same axons were delivering — read the flag and remove the process; C4d accumulates at the synapses in the target field and instructs their pruning through LilrB2; and Ephexin-5, elevated by amyloid and no longer degraded by a depleted EphB2, forbids the replacement of what was lost. Each of those steps is measured. The order in which this dissertation assembles them is not, and Section XIII says so.

Two therapeutic propositions follow, and they are unlike the propositions the deficiency model generates.

The first is an excitability brake rather than a transmitter replacement. If the coerulean axon is destroyed by the coerulean neuron's own firing, then the intervention is not to supply more norepinephrine or to stimulate the failing nucleus but to restore the control on its pacemaking — and *KCNIP4* is a named, human-validated, experimentally tractable candidate for what that control consists of, with a precedent in an adjacent wake-promoting nucleus where restoring a potassium brake is already proposed as a therapeutic strategy. The risk in this proposition is explicit and stated in Section XIV: quieting the locus coeruleus may preserve its arbor and impair the cognition that arbor supports, and the resolution of that tension is an empirical matter that has not been settled.

The second is an intact night. The locus coeruleus is asked to fall silent in REM and to oscillate slowly in NREM, and it is the only neuronal population in the brain granted a scheduled total off-period. On the mechanism established here, that off-period is the interval during which its axons are not being flagged for removal, and the NREM oscillation is the pump that clears the extracellular space of what damages it. Sleep is therefore not a lifestyle recommendation appended to a molecular argument. It is the tempo control on the clock that matters, and it is the only element of this entire chronology that is modifiable today, without a drug, in the decades during which the cell bodies are still all present and countable.

The locus coeruleus sickens in the third decade and dies in the eighth. The disease is decided in between — and it is decided at the far end of the axon, in the dark, while the count of cells is still normal.

References

Aston-Jones G, Bloom FE. Activity of norepinephrine-containing locus coeruleus neurons in behaving rats anticipates fluctuations in the sleep-waking cycle. *Journal of Neuroscience*. 1981;1(8):876–886.

Bondareff W, Mountjoy CQ, Roth M. Loss of neurons of origin of the adrenergic projection to cerebral cortex (nucleus locus ceruleus) in senile dementia. *Neurology*. 1982;32(2):164–168.

Braak H, Del Tredici K. The pathological process underlying Alzheimer's disease in individuals under thirty. *Acta Neuropathologica*. 2011;121(2):171–181.

Braak H, Thal DR, Ghebremedhin E, Del Tredici K. Stages of the pathologic process in Alzheimer disease: age categories from 1 to 100 years. *Journal of Neuropathology and Experimental Neurology*. 2011;70(11):960–969.

Brott BK, Raissi AJ, Micheva KD, Vielmetter J, Mendes MS, Baccus CJ, Huang J, Shatz CJ. C4d, a high-affinity LilrB2 ligand, is elevated in Alzheimer's disease and mediates synapse pruning. *Proceedings of the National Academy of Sciences USA*. 2025;122(38):e2519253122. PMID 40966293.

Bueichekú E, Diez I, Kim CM, Becker JA, Koops EA, Kwong K, Papp KV, Salat DH, Bennett DA, Rentz DM, Sperling RA, Johnson KA, Sepulcre J, Jacobs HIL. Spatiotemporal patterns of locus coeruleus integrity predict cortical tau and cognition. *Nature Aging*. 2024;4(5):625–637. PMID 38664576.

Chalermpalanupap T, Schroeder JP, Rorabaugh JM, Liles LC, Lah JJ, Levey AI, Weinshenker D. Locus coeruleus ablation exacerbates cognitive deficits, neuropathology, and lethality in P301S tau transgenic mice. *Journal of Neuroscience*. 2018;38(1):74–92. PMID 29133432.

Cissé M, Halabisky B, Harris J, Devidze N, Dubal DB, Sun B, Orr A, Lotz G, Kim DH, Hamto P, Ho K, Yu GQ, Mucke L. Reversing EphB2 depletion rescues cognitive functions in Alzheimer model. *Nature*. 2011;469(7328):47–52.

Crapser JD, Spangenberg EE, Barahona RA, Arreola MA, Hohsfield LA, Green KN. Microglia facilitate loss of perineuronal nets in the Alzheimer's disease brain. *EBioMedicine*. 2020;58:102919.

Dahl MJ, Mather M, Düzel S, Bodammer NC, Lindenberger U, Kühn S, Werkle-Bergner M. Rostral locus coeruleus integrity is associated with better memory performance in older adults. *Nature Human Behaviour*. 2019;3(11):1203–1214.

Dharshini SAP, Sanz-Ros J, Pan J, Tang W, Vallejo K, Otero-Garcia M, Cobos I. Molecular signatures of resilience to Alzheimer's disease in neocortical layer 4 neurons. *Nature Communications*. 2026. (Preprint: *bioRxiv* 2024.11.03.621787; PMC11580857.)

Ehrenberg AJ, Nguy AK, Theofilas P, Dunlop S, Suemoto CK, Di Lorenzo Alho AT, Leite RP, Diehl Rodriguez R, Mejia MB, Rüb U, Farfel JM, de Lucena Ferretti-Rebustini RE, Nascimento CF, Nitrini R, Pasquallucci CA, Jacob-Filho W, Miller B, Seeley WW, Heinsen H, Grinberg LT. Quantifying the accretion of hyperphosphorylated tau in the locus coeruleus and dorsal raphe nucleus: the pathological building blocks of early Alzheimer's disease. *Neuropathology and Applied Neurobiology*. 2017;43(5):393–408. PMID 28117917.

Elrod R, Peskind ER, DiGiacomo L, Brodkin KI, Veith RC, Raskind MA. Effects of Alzheimer's disease severity on cerebrospinal fluid norepinephrine concentration. *American Journal of Psychiatry*. 1997;154(1):25–30.

Ghosh A, Torraville SE, Mukherjee B, Walling SG, Martin GM, Harley CW, Yuan Q. An experimental model of Braak's pretangle proposal for the origin of Alzheimer's disease: the role of locus coeruleus in early symptom development. *Alzheimer's Research and Therapy*. 2019;11(1):59. PMID 31266535.

German DC, Manaye KF, White CL 3rd, Woodward DJ, McIntire DD, Smith WK, Kalaria RN, Mann DM. Disease-specific patterns of locus coeruleus cell loss. *Annals of Neurology*. 1992;32(5):667–676. PMID 1449247.

Grudzien A, Shaw P, Weintraub S, Bigio E, Mash DC, Mesulam MM. Locus coeruleus neurofibrillary degeneration in aging, mild cognitive impairment and early Alzheimer's disease. *Neurobiology of Aging*. 2007;28(3):327–335. PMID 16574280.

Hauglund NL, Andersen M, Tokarska K, Radovanovic T, Kjaerby C, Sørensen FL, Bojarowska Z, Untiet V, Ballestero SB, Kolmos MG, Weikop P, Hirase H, Nedergaard M. Norepinephrine-mediated slow vasomotion drives glymphatic clearance during sleep. *Cell*. 2025;188(3):606–622.e17.

Heneka MT, Nadrigny F, Regen T, Martinez-Hernandez A, Dumitrescu-Ozimek L, Terwel D, Jardanhazi-Kurutz D, Walter J, Kirchhoff F, Hanisch UK, Kummer MP. Locus ceruleus controls Alzheimer's disease pathology by modulating microglial functions through norepinephrine. *Proceedings of the National Academy of Sciences USA*. 2010;107(13):6058–6063.

Hong S, Beja-Glasser VF, Nfonoyim BM, Frouin A, Li S, Ramakrishnan S, Merry KM, Shi Q, Rosenthal A, Barres BA, Lemere CA, Selkoe DJ, Stevens B. Complement and microglia mediate early synapse loss in Alzheimer mouse models. *Science*. 2016;352(6286):712–716.

Jacobs HIL, Becker JA, Kwong K, Engels-Domínguez N, Prokopiou PC, Papp KV, Properzi M, Hampton OL, d'Oleire Uquillas F, Sanchez JS, Rentz DM, El Fakhri G, Normandin MD, Price JC, Bennett DA, Sperling RA, Johnson KA. In vivo and neuropathology data support locus coeruleus integrity as indicator of Alzheimer's disease pathology and cognitive decline. *Science Translational Medicine*. 2021;13(612):eabj2511.

Kang SS, Liu X, Ahn EH, Xiang J, Manfredsson FP, Yang X, Luo HR, Liles LC, Weinshenker D, Ye K. Norepinephrine metabolite DOPEGAL activates AEP and pathological Tau aggregation in locus coeruleus. *Journal of Clinical Investigation*. 2020;130(1):422–437.

Kelly SC, He B, Perez SE, Ginsberg SD, Mufson EJ, Counts SE. Locus coeruleus cellular and molecular pathology during the progression of Alzheimer's disease. *Acta Neuropathologica Communications*. 2017;5(1):8. PMID 28109312.

Kim T, Vidal GS, Djurisic M, William CM, Birnbaum ME, Garcia KC, Hyman BT, Shatz CJ. Human LilrB2 is a beta-amyloid receptor and its murine homolog PirB regulates synaptic plasticity in an Alzheimer's model. *Science*. 2013;341(6152):1399–1404.

Kobro-Flatmoen A, Nagelhus A, Witter MP. Reelin-immunoreactive neurons in entorhinal cortex layer II selectively express intracellular amyloid in early Alzheimer's disease. *Neurobiology of Disease*. 2016;93:172–183. PMID 27195475.

Lester E, Ooi FK, Bakkar N, Ayers J, Woerman AL, Wheeler J, Bowser R, Carlson GA, Prusiner SB, Parker R. Tau aggregates are RNA-protein assemblies that mislocalize multiple nuclear speckle components. *Neuron*. 2021;109(10):1675–1691.e9. PMID 33848474.

Li SB, Damonte VM, Chen C, Wang GX, Kebschull JM, Yamaguchi H, Bian WJ, Purmann C, Pattni R, Urban AE, Mourrain P, Kauer JA, Scherrer G, de Lecea L. Hyperexcitable arousal circuits drive sleep instability during aging. *Science*. 2022;375(6583):eabh3021.

Margolis SS, Salogiannis J, Lipton DM, Mandel-Brehm C, Wills ZP, Mardinly AR, Hu L, Greer PL, Bikoff JB, Ho HY, Soskis MJ, Sahin M, Greenberg ME. EphB-mediated degradation of the RhoA GEF Ephexin5 relieves a developmental brake on excitatory synapse formation. *Cell*. 2010;143(3):442–455. PMID 21029865.

Matchett BJ, Grinberg LT, Theofilas P, Murray ME. The mechanistic link between selective vulnerability of the locus coeruleus and neurodegeneration in Alzheimer's disease. *Acta Neuropathologica*. 2021;141(5):631–650.

Mather M, Harley CW. The locus coeruleus: essential for maintaining cognitive function and the aging brain. *Trends in Cognitive Sciences*. 2016;20(3):214–226.

Meyer C, Niedermeier T, Feyen PLC, Strübing FL, Rauchmann BS, Karali K, Gentz J, Tillmann YE, Landgraf NF, Rumpf SL, Ochs K, Wind-Mark K, Biechele G, Wagner J, Guersel S, Kurz CI, Schweiger M, Prtvar D, Shi Y, Banati RB, Liu GJ, Middleton RJ, Mitteregger-Kretzschmar G, Perneczky R, Koeglsperger T, Neher JJ, Tahirovic S, Brendel M, Herms J, Paeger L. Early locus coeruleus noradrenergic axon loss drives olfactory dysfunction in Alzheimer's disease. *Nature Communications*. 2025;16:7338. PMID 40781079.

Morawski M, Brückner G, Jäger C, Seeger G, Arendt T. Neurons associated with aggrecan-based perineuronal nets are protected against tau pathology in subcortical regions in Alzheimer's disease. *Neuroscience*. 2010;169(3):1347–1363.

Morohashi Y, Hatano N, Ohya S, Takikawa R, Watabiki T, Takasugi N, Imaizumi Y, Tomita T, Iwatsubo T. Molecular cloning and characterization of CALP/KChIP4, a novel EF-hand protein interacting with presenilin 2 and voltage-gated potassium channel subunit Kv4. *Journal of Biological Chemistry*. 2002;277(17):14965–14975. PMID 11847232.

Oh J, Eser RA, Ehrenberg AJ, Morales D, Petersen C, Kudlacek J, Dunlop SR, Theofilas P, Resende EPF, Cosme C, Alho EJL, Spina S, Walsh CM, Miller BL, Seeley WW, Bittencourt JC, Neylan TC, Heinsen H, Grinberg LT. Profound degeneration of wake-promoting neurons in Alzheimer's disease. *Alzheimer's and Dementia*. 2019;15(10):1253–1263. PMID 31416793.

Rorabaugh JM, Chalermpalanupap T, Botz-Zapp CA, Fu VM, Lembeck NA, Cohen RM, Weinshenker D. Chemogenetic locus coeruleus activation restores reversal learning in a rat model of Alzheimer's disease. *Brain*. 2017;140(11):3023–3038. PMID 29053824.

Sekar A, Bialas AR, de Rivera H, Davis A, Hammond TR, Kamitaki N, Tooley K, Presumey J, Baum M, Van Doren V, Genovese G, Rose SA, Handsaker RE; Schizophrenia Working Group of the Psychiatric Genomics Consortium; Daly MJ, Carroll MC, Stevens B, McCarroll SA. Schizophrenia risk from complex variation of complement component 4. *Nature*. 2016;530(7589):177–183.

Sell GL, Schaffer TB, Margolis SS. Reducing expression of synapse-restricting protein Ephexin5 ameliorates Alzheimer's-like impairment in mice. *Journal of Clinical Investigation*. 2017;127(5):1963–1968. PMID 28346227.

Szot P, White SS, Greenup JL, Leverenz JB, Peskind ER, Raskind MA. Compensatory changes in the noradrenergic nervous system in the locus ceruleus and hippocampus of postmortem subjects with Alzheimer's disease and dementia with Lewy bodies. *Journal of Neuroscience*. 2006;26(2):467–478.

Theofilas P, Ehrenberg AJ, Dunlop S, Di Lorenzo Alho AT, Nguy A, Leite REP, Rodriguez RD, Mejia MB, Suemoto CK, Ferretti-Rebustini REL, Polichiso L, Nascimento CF, Seeley WW, Nitrini R, Pasqualucci CA, Jacob Filho W, Rueb U, Neuhaus J, Heinsen H, Grinberg LT. Locus coeruleus volume and cell population changes during Alzheimer's disease progression: a stereological study in human postmortem brains with potential implication for early-stage biomarker discovery. *Alzheimer's and Dementia*. 2017;13(3):236–246. PMID 27513978.

Van Egroo M, van Hooren RWE, Jacobs HIL. Associations between locus coeruleus integrity and nocturnal awakenings in the context of Alzheimer's disease plasma biomarkers: a 7T MRI study. *Alzheimer's Research and Therapy*. 2021;13(1):159.

Weinshenker D. Long road to ruin: noradrenergic dysfunction in neurodegenerative disease. *Trends in Neurosciences*. 2018;41(4):211–223.

Wilson RS, Nag S, Boyle PA, Hizel LP, Yu L, Buchman AS, Schneider JA, Bennett DA. Neural reserve, neuronal density in the locus ceruleus, and cognitive decline. *Neurology*. 2013;80(13):1202–1208.

Xie L, Kang H, Xu Q, Chen MJ, Liao Y, Thiyagarajan M, O'Donnell J, Christensen DJ, Nicholson C, Iliff JJ, Takano T, Deane R, Nedergaard M. Sleep drives metabolite clearance from the adult brain. *Science*. 2013;342(6156):373–377.

Zarow C, Lyness SA, Mortimer JA, Chui HC. Neuronal loss is greater in the locus coeruleus than nucleus basalis and substantia nigra in Alzheimer and Parkinson diseases. *Archives of Neurology*. 2003;60(3):337–341. PMID 12633144.